

Classification and Reconstruction of Surfaces from Point Clouds of Man-made Objects

Falko Schindler Wolfgang Förstner
Department of Photogrammetry
University of Bonn
Nussallee 15, 53115 Bonn, Germany

<http://www.ipb.uni-bonn.de/{schindler,foerstner}>

Jan-Michael Frahm
Department of Computer Science
University of North Carolina
Chapel Hill, NC, 27599, USA

<http://www.cs.unc.edu/~jmf>

Abstract

We present a novel surface model and reconstruction method for man-made environments that take prior knowledge about topology and geometry into account. The model favors but is not limited to horizontal and vertical planes that are pairwise orthogonal. The reconstruction method does not require one particular class of sensors, as long as a triangulated point cloud is available. It delivers a complete 3D segmentation, parametrization and classification for both surface regions and inter-plane relations. By working on a pre-segmentation we reduce the computational cost and increase robustness to noise and outliers. All reasoning is statistically motivated, based on a few decision variables with meaningful interpretation in measurement space. We demonstrate our reconstruction method for visual reconstructions and laser range data.

1. Introduction

Despite the tremendous progress in computer vision we still lack the ability of a full semantic understanding of the surrounding environment comparable to the human understanding. This interpretation of the environment is key, e. g., for applications in robotics for navigational purposes. Recent research in computer vision [4, 5, 7, 8] started to address the problem by enforcing planar surfaces within reconstructions of urban environments and indoor scenes. Rather than limiting the process of enforcing planar reconstructions to one particular class of sensors, in this paper we propose a method that enforces the prior knowledge on any input mesh provided. Specifically our method finds all surface regions, which can be described by a low dimensional parameterized surface model, e. g., a plane.

We assume a triangulated point cloud to be available either derived from image collections or videos using structure-from-motion techniques or directly from laser

range data. We wish to find uniformly describable surface regions as well as their optimal parametrization. Moreover, we are interested in a complete labeling of the mesh, rather than detecting dominant regions only, yielding undefined surface parts near boundaries and holes. Given the inherent constraints of man-made environments we favor special surface classes and relations between neighboring regions, e. g., orthogonality of a horizontal and a vertical plane.

Our contribution is a surface model for man-made scenes that takes prior knowledge about geometry and topology into account. It incorporates classes for special regions like orthogonal and vertical planes as well as relations between neighboring regions like identity and orthogonality. Furthermore we present a reconstruction method for applying the surface model to triangulated meshes delivering both a globally optimal surface classification and parametrization. By working on a pre-segmented surface we reduce the computational cost while increasing robustness to noise and outliers.

Our proposed surface model naturally encodes world knowledge in an intuitive way. Therefore it is easily extendible with more classes and relations. All reasoning during the classification and parameter estimation is statistically motivated and based on the statistics of the measurements used. Hence our decision variables have a meaningful interpretation in the measurement space. This is an advantage as all parameters can be chosen based on commonly known uncertainties, e. g., the surface accuracy.

The remainder of the paper is organized as follows. First we discuss related work in Section 2, before describing the proposed surface model in detail in Section 3 and introducing our novel reconstruction method in Section 4. In Section 5 we demonstrate the performance of our approach on both stereo image reconstructions and laser range data.

2. Related work

As argued in [2], extensive work exists on segmentation of depth and range images, not being applicable to our problem due to its restriction to 2.5D. Others address reconstruction tasks in computer-aided design and thus solely focus on laser scans or synthetic data with a very small amount of noise and rare outliers compared to typical camera based reconstruction [2].

RANSAC-driven approaches, e. g. [12, 8], due to computational complexity are often limited to detecting the most dominant surface segments only. In general they do not yield watertight segmentations as they do not model the inter-plane relations and hence fail to enforce intersections of nearby planes, which is an intrinsic property in man-made environments. Moreover an extension to more complex surface models than planes involves a rapidly increasing computational complexity.

Plane-sweeping techniques and their derivatives, e. g. [13, 7], are only using local planar surface approximations and hence fail to obtain larger scale scene planes. In [7] the authors obtain likely surface normal directions to improve the local plane approximations for a gain in accuracy and formulating a location prior. In contrast to our proposed model this knowledge is not deployed to obtain any global surface model.

Manhattan-world stereo, e. g. [4, 5], is a scene model consisting of a set of strictly orthogonal planes. This fails to model any general plane like roof tops as well as to explain any other scene geometry present as pointed out in [8]. Our proposed method is not limited to a strictly orthogonal plane model. In [5] the fusion of multiple depth maps into an orthogonal plane world model involves a volumetric world model that gets computational infeasible with growing size.

Mesh segmentation, e. g. [1], is normally applied to laser scanning data and fails on 3D points reconstructed from multi-view stereo typically featuring higher noise levels than laser scanning. The very promising approach [1] can handle planes as well as higher order surfaces, but does not involve constraints, like, e. g., orthogonality between neighboring planes, and is not devised for the noisier stereo reconstruction data. Only local decisions are made that could lead into problems, when noise and outliers are present.

Region growing is widely used in segmentation tasks, e. g., the authors of [2] present a sophisticated approach for piecewise planar 3D reconstruction from point clouds and show impressive results. Unfortunately their complex algorithm is not extendible for any other architectural surface types, which typically occur in man-made environments, e. g., cylindrical surfaces modeling columns. Furthermore, they exploits visibility constraints that might not be available for every 3D point cloud, depending on the acquisition technique and post-processing.

Image-based clustering approaches like [16] yield a high quality 2D separation of the different observed surfaces but do not provide a 3D interpretation of the surface, which is essential to enforce the inter-plane relations in urban environments to ensure water tight surfaces. This modeling of inter-plane relations is a key advantage of our proposed method.

3. Surface model

We propose a model for describing a triangulated surface consisting of piecewise analytically describable segments. Different semantic areas, in the following called regions, can contain multiple segments that are expected to have identical surface parameters. We assume that there is a dependency between the parametrization of neighboring surface segments, their class label and the explicit relation between them.

3.1. Model for global classification

Each of the K surface segments \mathcal{S}_k with $k = 1 \dots K$ is observed in the form of Euclidean coordinate vectors of the 3D points $\mathcal{X}_k = \{\mathbf{x}_{k,i}\}$ with $i = 1 \dots I_k$. Our surface model connects the parametrization θ_k with the class labeling c_k for each surface segment and the label $r_n = r_n(k, k')$ for the relation n between two neighboring segments.

Figure 1 shows the relation n between two segments k and k' as a factor graph, encoding the unary and binary potentials ϕ_k and ϕ_n

$$\phi_k(c_k, \theta_k) = P(\theta_k | c_k)P(c_k), \quad (1)$$

$$\phi_n(c_k, c_{k'}, r_n, \theta_k, \theta_{k'}) = P(\theta_k, \theta_{k'} | r_n)P(c_k, c_{k'} | r_n)P(r_n). \quad (2)$$

The consistency term $P(c_k, c_{k'} | r_n)$ links neighboring classes and their relation. The joint probability of the configuration is given by

$$P(c, r, \theta) \propto \prod_k \phi_k(c_k, \theta_k) \prod_n \phi_n(c_k, c_{k'}, r_n, \theta_k, \theta_{k'}). \quad (3)$$

Next we detail the proposed parametrization and a set of classes and relations, suited for man-made scenes.

3.2. Parametrization

In the most general case we represent each surface segment \mathcal{S}_k by an implicit surface function f and a parametrization θ_k satisfying

$$f(\mathbf{x}_{k,i}, \theta_k) = 0 \quad (4)$$

for each 3D point $\mathbf{x}_{k,i}$ living on segment \mathcal{S}_k . A possible parametrization of f is the quadratic form $f(\mathbf{x}_{k,i}, \mathbf{Q}_k) = \mathbf{x}_{k,i}^T \mathbf{Q}_k \mathbf{x}_{k,i}$ with \mathbf{Q}_k being a 4×4 symmetric, homogeneous parameter matrix of a quadric, i. e. a second order surface.

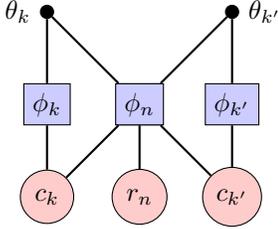


Figure 1. Factor graph modeling the class labeling c_k and $c_{k'}$ of two surface segments k and k' as well as the relation r_n between both of them, all depending on the surface parametrization θ_k and $\theta_{k'}$.

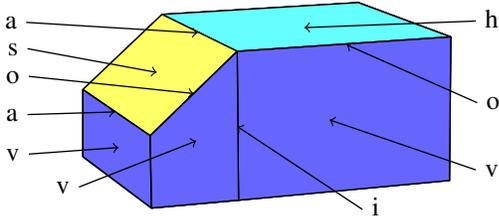


Figure 2. Example instance of our surface model incorporating *slanted* (s), *vertical* (v) and *horizontal* (h) planes as well as relations *arbitrary* (a), *orthogonal* (o) and *identical* (i) between two of them.

Note that this representation is linear in the parameters \mathbf{Q}_k , despite of a curved surface shape.

Since many man-made objects are piecewise planar, we will use a more restrictive model in the following. We model first order surfaces, i. e. planes, parametrized with a 4-vector $\theta = [\mathbf{n}^\top, -d]^\top$ with the normalized normal vector \mathbf{n} and the distance d of the plane to the origin. Ideally all points $\mathbf{x}_{k,i}$ living on a segment \mathcal{S}_k fulfill the condition

$$f(\mathbf{x}_{k,i}, \theta_k) = \mathbf{x}_{k,i}^\top \theta_k = 0. \quad (5)$$

3.3. Surface classes

Besides arbitrarily oriented planes, in the following called *slanted*, we often observe *horizontal* and *vertical* planes (Table 1). This requires our coordinate system to be aligned to the scene’s vertical direction, similar to [7].

While a *slanted* plane has all three degrees of freedom, the other two classes involve constraints to the normal vector \mathbf{n} . The expected value for one or two of its components is zero. Instead of enforcing crisp constraints we treat them as prior. E. g. the prior for the z -component of the normal for a vertical plane has prior $n_z \sim N(0, \sigma_{n_z}^2)$ with $\sigma_{n_z}^2 = \sigma_\alpha^2$ since $n_z = \sin \alpha$. The angle σ_α is the expected deviation from the ideal model, being 0.25° in all our experiments.

Different a-priori probabilities enforce vertical and horizontal planes despite the constraints to be fulfilled. In general, the probability for a segment being *slanted* η_s is small

and the other two priors are of equal size $(1 - \eta_s)/2$, so that all priors sum to one. In all our experiments we use $\eta_s = 1\%$, yielding most reasonable results.

Next we detail our proposed method to model the relationships in between different surface primitives.

Class	Parameter distribution	Prior
<i>slanted</i>	—	η_s
<i>vertical</i>	$n_z \sim N(0, \sigma_\alpha^2)$	$(1 - \eta_s)/2$
<i>horizontal</i>	$n_x, n_y \sim N(0, \sigma_\alpha^2)$	$(1 - \eta_s)/2$

Table 1. Classes for separate surface segments. We model planes with different orientation, fixing up to two degrees of freedom. In order to penalize using the class *slanted* without any constraint we choose a small a-priori probability η_s .

3.4. Inter-primitive relations

Two neighboring planes can intersect orthogonally or arbitrarily, suggesting two relations *orthogonal* and *arbitrary*. However, when modeling neighboring surface segments, they not necessarily have to differ. Therefore we add the identity relation (Table 2).

Again we model the relations using their priors. In case of orthogonal normal vectors the dot product has to be zero, while for identical normal vectors the cross product vanishes. Furthermore, the center of gravity $\bar{\mathbf{x}}_k$ of all points $\mathbf{x}_{k,i}$ on segment \mathcal{S}_k has to lie within the neighboring plane $\theta_{k'}$, enforcing *identical* planes not only to be parallel but coincident. The variances of these constraints are derived in a similar way like in Section 3.3 from a deviation angle σ_α and a tolerance σ_δ for the distance between two identical planes.

Table 2 also lists multiple pairs of surface classes for each relation, since specific relations are only allowed for certain classes. E. g. the identity relation can only occur between two segments of the same class. The permission and prohibition of certain combinations are encoded in the consistency term in (2).

Relations *orthogonal* and *identical* have equal a-priori probability, while we chose a small value $\eta_a = 1\%$ for arbitrarily related segments.

Figure 2 shows all modelled surface classes and relations.

After defining the surface parameterization and determining the modeling of their pairwise relations we next discuss our proposed reconstruction method.

4. Reconstruction method

Given a triangulated point cloud, i. e. a surface mesh, we first apply a pre-segmentation to obtain surface segments. Then we need to optimize the model w. r. t. the joint probability (3). We split this task into two parts, first the inference of globally optimal classes and relations given the

Relation	Parameter distribution	Classes	Prior
<i>arbitrary</i>	—	<i>hs, vs, vv, ss</i>	η_a
<i>orthogonal</i>	$\mathbf{n}_k^\top \mathbf{n}_{k'} \sim N(0, 2\sigma_\alpha^2)$	<i>vh, vv, vs, ss</i>	$(1 - \eta_a)/2$
<i>identical</i>	$(\bar{\mathbf{x}}_k^\top \mathbf{n}_{k'} - d_{k'}) + (\bar{\mathbf{x}}_{k'}^\top \mathbf{n}_k - d_k) \sim N(0, 2\sigma_\delta^2),$ $\mathbf{n}_k \times \mathbf{n}_{k'} \sim N(\mathbf{0}, \sigma_\alpha^2 I_3)$	<i>vv, hh, ss</i>	$(1 - \eta_a)/2$

Table 2. Relations between neighboring surface segments. In case of *orthogonal* planes we fix the angle between both normal vectors, while *identical* planes have to coincide in orientation and location. Relations can occur in between of certain pairs of vertical, horizontal and slanted planes. In order to penalize using the relation *arbitrary* without any constraint we choose a small a-priori probability η_a .

initial surface parameters, and second the estimation of a globally optimal parametrization using the labeling as soft constraints. Finally we merge identical segments to larger, semantically interpretable surface regions and compute a final parameter estimation for these regions.

4.1. Pre-segmentation

Starting from a triangulated surface the smallest accessible surface elements are triangles. When, however, the surface is very rough, as it is often the case with visual reconstructions, a single triangles information can be significantly disturbed and hence its properties can be misleading.

Therefore we pre-segment the given surface using Fast Marching Farthest Point Sampling [10], grouping triangles based on their geodesic distance as well as their local curvature [11]. Finding seed points is part of the algorithm and does not need any additional computation. The distance function can be a combination of displacement, difference between normals or even radiometric features. As shown in [15], Fast Marching can be parallelized efficiently, enabling to process point clouds with millions of points in far less than a second.

4.2. Initial segment parameter estimation

We estimate the parameters $\hat{\theta}_k$ for each segment \mathcal{S}_k individually from the observed 3D points \mathcal{X}_k , minimizing

$$\Omega_k = \sum_i (\hat{\mathbf{x}}_{k,i} - \mathbf{x}_{k,i})^\top \Sigma_{x_k, i x_k, i}^{-1} (\hat{\mathbf{x}}_{k,i} - \mathbf{x}_{k,i}) \quad (6)$$

subject to the functional model (5) and constraints h_k between unknown parameters θ_k , e. g., restricting the normal vector to unit length $|\mathbf{n}_k| = 1$ in case of planes.

After determining the initial parametrization we next describe our proposed global parameter optimization.

4.3. Globally optimal segment labeling

First we compute probabilistic measures for each surface segment \mathcal{S}_k belonging to a model c_k represented by a set of constraints \mathcal{H}_{c_k} . In the case of planar surfaces the constraint

$$\mathbf{h}_{c_k} = h_{c_k}(\hat{\theta}_k) = \mathbf{0} \quad (7)$$

can be written in form of $\mathbf{h}_{c_k} = H_{c_k}^\top \hat{\theta}_k$ with

$$H_{c_k}^\top = \begin{bmatrix} 0 & 0 & 1 & 0 \end{bmatrix} \quad (8)$$

for *vertical* planes and

$$H_{c_k}^\top = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \quad (9)$$

for *horizontal* planes. For our experiments we obtain the covariance of the constraints, Σ_{hh} , via variance propagation from Σ_{aa} . Alternatively we could use the models uncertainty formulated in Tables 1 and 2.

In case the constraint (7) is fulfilled the weighted sum $\Omega_{c_k} = \mathbf{h}_{c_k}^\top \Sigma_{hh, c_k}^{-1} \mathbf{h}_{c_k}$ is assumed to be χ^2 -distributed with $|\mathcal{H}_{c_k}|$ degrees of freedom. We obtain the likelihood $p(\theta_k | c_k)$ from the χ^2 -density function $p(\theta_k | c_k) = p(\Omega_{c_k})$. Similarly we compute probabilities $P(\theta_k, \theta_{k'} | r_n)$ for the relations r_n between two neighboring segments k and k' .

We infer an optimal labeling c_k and r_n using the sum-product algorithm maximizing the joint probability (3). The sum-product algorithm requires the graph to be tree-structured for guaranteeing exact inference, which is not provided by our graph. Nevertheless empirically we determined that we find a close approximation of the optimal labeling as well as for the marginal probabilities $P(c_k | \theta)$ and $P(r_n | \theta)$.

4.4. Constraint segment parameter estimation

We similarly estimate the parameters θ_k from the observed 3D points \mathcal{X}_k as described Section 4.2 with the added soft constraints formulated in Tables 1 and 2.

For each segment k we introduce class-dependent constraints h_{c_k} for all classes c_k weighted with the corresponding probability $P(c_k | \theta)$. For each relation n we introduce relation-dependent constraints h_{r_n} for all relations r_n weighted with the corresponding probability $P(r_n | \theta)$. Therefore, we minimize the energy

$$\begin{aligned} \Omega = & \sum_k \sum_i (\hat{\mathbf{x}}_{k,i} - \mathbf{x}_{k,i})^\top \Sigma_{x_k, i x_k, i}^{-1} (\hat{\mathbf{x}}_{k,i} - \mathbf{x}_{k,i}) \\ & + \sum_k \sum_c P(c_k | \theta) \mathbf{h}_{c_k}^\top \Sigma_{hh, c}^{-1} \mathbf{h}_{c_k} \\ & + \sum_n \sum_r P(r_n | \theta) \mathbf{h}_{r_n}^\top \Sigma_{hh, r}^{-1} \mathbf{h}_{r_n}. \end{aligned} \quad (10)$$

subject to the functional model (5) and constraints h_k as in Section 4.2. Here we neglect the mutual dependencies between segments and relations.

4.5. Merging segments to regions

In order to improve the segmentation from the very beginning, we merge all segments being most likely *identical*. We name the elements of the new segmentation *regions*.

Assigning common region numbers to all segments linked by the relation *identical* can be formulated as graph coloring task. Our segments are graph nodes and links between two of them are set if they are *identical*. Now we label each connected component with a unique region label. This can be done applying a Dulmage-Mendelsohn decomposition [3] for the adjacency matrix of the graph.

4.6. Constraint region parameter estimation

Finally, we estimate parameters for all regions from Section 4.5. The class membership of a region and the relation between regions is not unique. We simply apply constraints from the most frequent class within each region and from the most frequent relation between two regions.

5. Experiments

We demonstrate the proposed reconstruction method using our surface model with various data sets arising from different data acquisition techniques. After describing the generation of each data set in brief (Table 3), we illustrate the single steps and intermediate results of our proposed reconstruction method using the example of the CHURCH data set. Finally we present and discuss results obtained with the other data sets.

Name	Source	# points	# triangles
CHURCH	100 turntable images	16 277	32 192
CUBE	synthetically sampled	6 000	11 996
BOX	55 turntable images	26 516	52 393
BLOCK	close range laser scan	15 492	30 382
HOUSE	close range laser scan	19 915	39 662
HALL	terrestrial laser scan	21 730	43 084
BLOCKCITY	48 turntable images	35 651	70 382
BRICKS	72 turntable images	39 480	78 503
HILL	structured light scanning	34 785	68 938

Table 3. All data sets used in our experiments.

5.1. Data sets

CUBE is a meshed point cloud generated synthetically from 6000 points sampled on a cube. It allows to control the amount of noise precisely and to check, whether our reconstruction yields correct decisions when knowing the true variance of the data points.

CHURCH, BOX, BLOCKCITY and BRICKS are images of objects mounted on a rotating turntable and captured using a static camera. We reconstructed the camera orientation using the structure-from-motion software Bundler [14] and reconstructed a dense point cloud using the Patch-based

Multi-view Stereo software (PMVS, [6]). We generated a mesh using the Poisson surface reconstruction [9], rather than computing a triangulation of the raw points. Although Poisson surfaces tend to smooth edges, which would be unfavorable to our purpose, we obtain useful results with an octree depth of 7 or 8.

BLOCK and HOUSE are triangulated high-resolution point-clouds of a small building model and a wooden block captured using a Perceptron ScanWorks V5 laser scanner mounted on a Romer measuring arm. This data set is characterized by very small noise in the order of tens of micrometers.

HALL is a terrestrial laser scan acquired from 7 stand points located around a small house. We generated a mesh by triangulating a randomly sampled subset of points.

HILL is a triangular mesh of a historic city model of Hamburg captured using structured light scanning.

5.2. The reconstruction method in detail

Figure 3 illustrates the intermediate steps of our reconstruction method. In this case we reconstruct a dense, meshed point cloud via structure-from-motion from 100 turntable images (3(a)). The resulting mesh is shown in Figure 3(b). We pre-segment the surface mesh using Fast Marching Farthest Point Sampling. Around 60 segments yield a good pre-segmentation, i. e., most importantly, all edges are preserved. On the other hand the segmentation is not too fine, but equally large and stable. We compute locally optimal parameters for each segment (3(d)) used for a first labeling of segments and relations (3(e)). Note the yellow segments within the blue and cyan, i. e. vertical and horizontal surfaces. In most cases such misclassified segments vanish after improving the labeling by incorporating binary relations between neighboring segments (3(f)). Some segments remain classified as *slanted*. E. g. the front is divided into two regions, since the original mesh is quite uneven compared to the noise. The determined classes and relations are used to constrain the parameter estimation and to merge *identical* segments to larger, semantically motivated regions shown in Figure 3(g). After optimizing the parametrization once more, now for surface regions rather than segments, we obtain an almost perfect, accurate reconstruction in Figure 3(h).

5.3. Results on various data sets

Figure 4 shows reconstruction results we obtained with various other data sets. A challenging data set is HALL due to its windows, doors and dormers that are not preserved in this experiment. The low point density on the roof does not allow reconstructing such details, in this case. The parametrization and classification, however, yields a clean generalization towards walls and roof, due to the prior towards verticality and orthogonality.

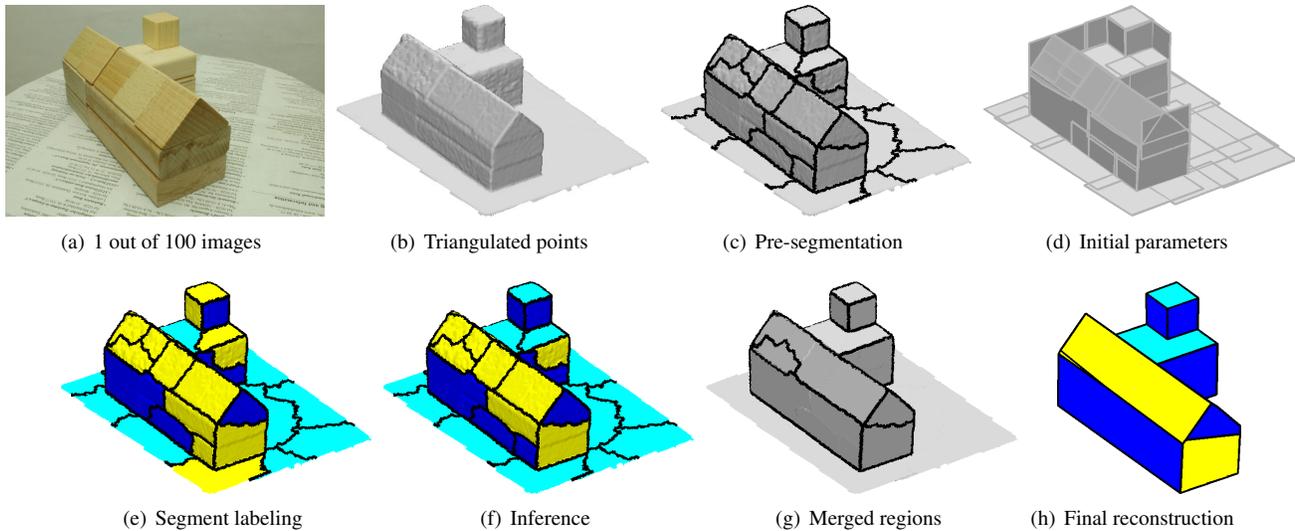


Figure 3. Reconstruction steps demonstrated using the example of the CHURCH data set. Classes in (e) and (f) are *vertical* (blue), *horizontal* (cyan) and *slanted* (yellow).

BLOCKCITY, BRICKS and HILL are not reconstructed perfectly. The compression is still large, while preserving the geometric structure accurately. Most of the observable over-segmentation is caused by inaccurately built objects, e. g., vertical wall being not exactly orthogonal to horizontal planes or slightly curved surface regions. This over-segmentation could be overcome by tuning the model uncertainties σ_α and σ_δ for each data set individually.

6. Conclusion

We presented a novel surface model for man-made scenes incorporating prior knowledge about geometry and topology. It favors vertical and horizontal planes, being pairwise orthogonal or identical, while arbitrarily oriented planes are modeled as well. Our proposed reconstruction method links surface parameter estimation and labeling, delivering a complete, coherent model. The surface model can easily be extended with other classes and relations. All control parameters have a meaningful interpretation.

We demonstrated our approach with triangulated point clouds from multi-view stereo and laser scans, obtaining satisfactory results.

References

- [1] M. Attene, B. Falcidieno, and M. Spagnuolo. Hierarchical Mesh Segmentation Based on Fitting Primitives. *VISUAL COMPUT*, 22(3):181–193, 2006. 2
- [2] A.-L. Chauve, P. Labatut, and J.-P. Pons. Robust Piecewise-planar 3D Reconstruction and Completion from Large-scale Unstructured Point Data. In *CVPR*, 2010. 2
- [3] A. Dulmage and N. Mendelsohn. Coverings of Bipartite Graphs. *Canadian Journal of Mathematics*, (10):517–534, 1958. 5
- [4] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski. Manhattan-world Stereo. In *CVPR*, pages 1422–1429. Ieee, June 2009. 1, 2
- [5] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski. Reconstructing Building Interiors from Images. In *ICCV*, pages 80–87, 2009. 1, 2
- [6] Y. Furukawa and J. Ponce. Accurate, Dense, and Robust Multi-view Stereopsis. In *CVPR*, pages 1362–1376, 2007. 5
- [7] D. Gallup, J.-M. Frahm, P. Mordohai, Q. Yang, and M. Pollefeys. Real-time Plane-sweeping Stereo with Multiple Sweeping Directions. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. 1, 2, 3
- [8] D. Gallup, J.-M. Frahm, and M. Pollefeys. Piecewise Planar and Non-planar Stereo for Urban Scene Reconstruction. In *CVPR*, pages 1418–1425. IEEE, 2010. 1, 2
- [9] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson Surface Reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, pages 61–70, New York, New York, USA, 2006. Eurographics Association. 5
- [10] C. Moenning and N. A. Dodgson. Fast Marching Farthest Point Sampling for Point Clouds and Implicit Surfaces, 2003. 4
- [11] F. Schindler and W. Förstner. Fast Marching for Robust Surface Segmentation. In *Photogrammetric Image Analysis, Lecture Notes in Computer Science*, 2011. 4
- [12] R. Schnabel, R. Wahl, and R. Klein. Efficient RANSAC for Point-Cloud Shape Detection. *Computer Graphics Forum*, 26(2):214–226, June 2007. 2
- [13] S. N. Sinha, D. Steedly, and R. Szeliski. Piecewise Planar Stereo for Image-based Rendering. In *ICCV*, 2009. 2

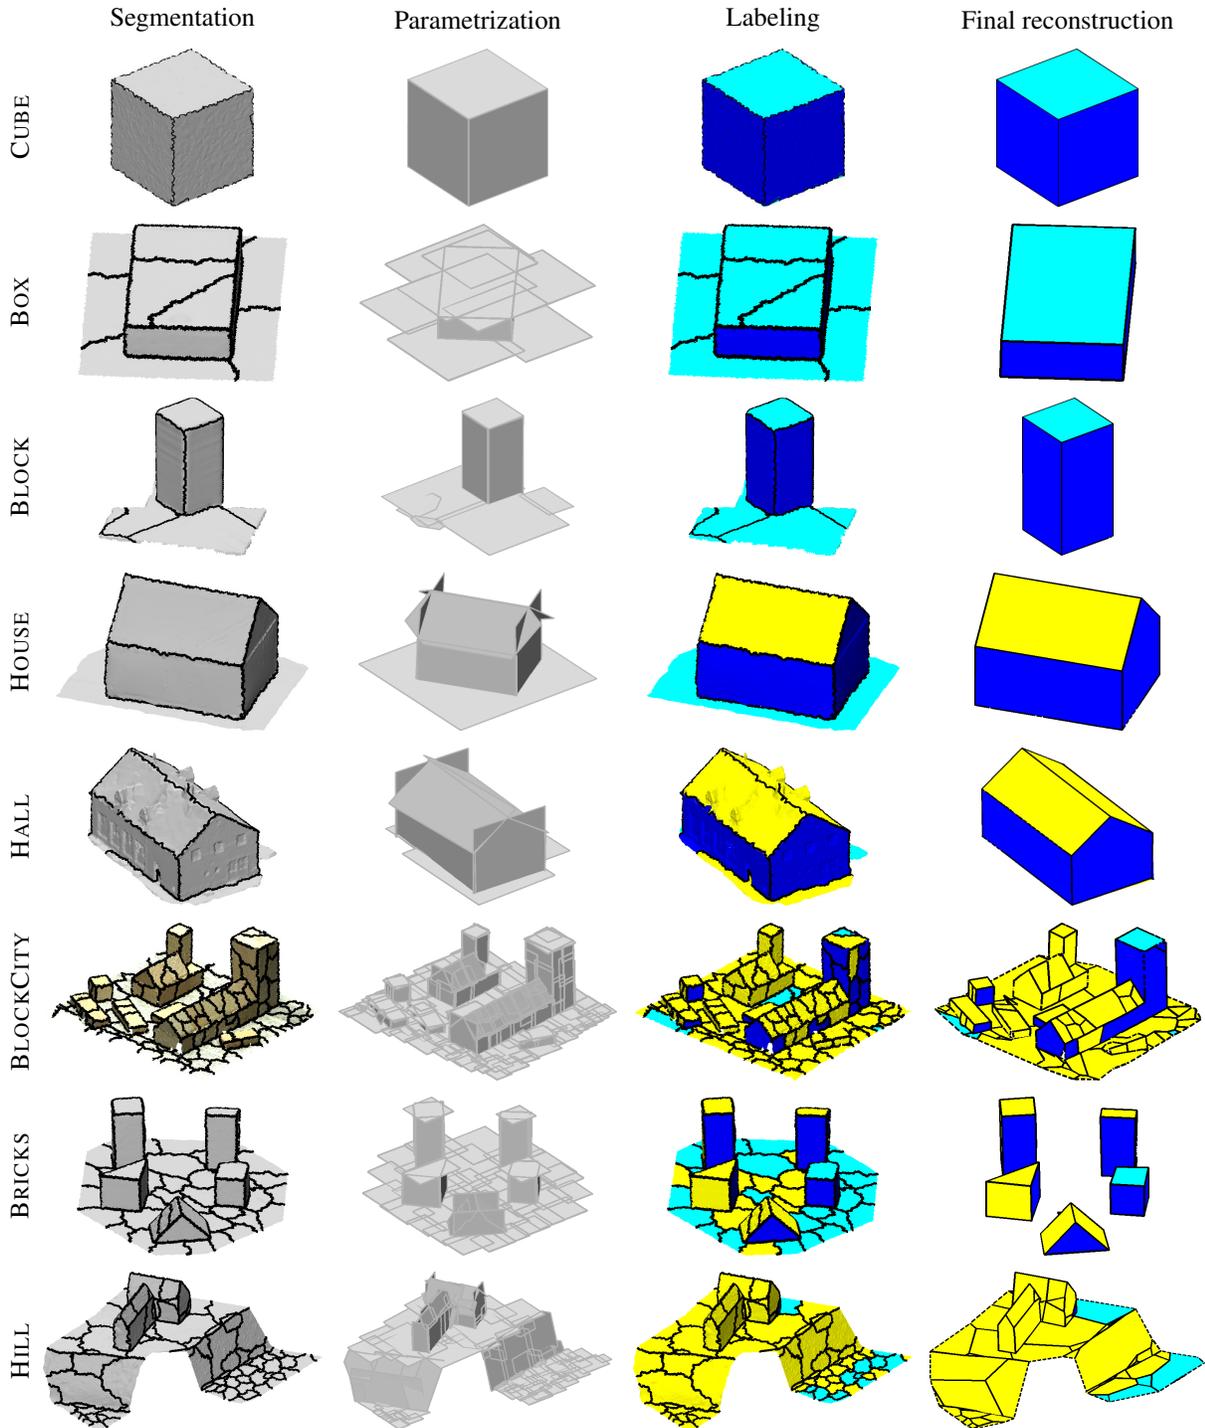


Figure 4. Segmentation, parametrization, labeling and final reconstruction for eight data sets.

[14] N. Snavely, S. M. Seitz, and R. Szeliski. Photo Tourism: Exploring Photo Collections in 3D. In *ACM T GRAPHIC*, pages 835–846, 2006. 5

[15] O. Weber, Y. S. Devir, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Parallel Algorithms for Approximation of

Distance Maps on Parametric Surfaces. *ACM Transactions on Graphics*, 27(4):1–16, Oct. 2008. 4

[16] G. Yu, M. Grossberg, G. Wolberg, and I. Stamos. Think Globally, Cluster Locally: A Unified Framework for Range Segmentation. In *Proc. 3DPVT*. Citeseer, 2008. 2